

Inference with Discriminative Posterior

Jarkko Salojärvi[†] Kai Puolamäki[‡] Eerika Savia[†]

Samuel Kaski[†]

Helsinki Institute for Information Technology
Department of Information and Computer Science
Helsinki University of Technology
P.O. Box 5400, FI-02015 TKK, Finland

[†] Author belongs to the Finnish Centre of Excellence in
Adaptive Informatics Research.

[‡] Author belongs to the Finnish Centre of Excellence in
Algorithmic Data Analysis Research.

Abstract

We study Bayesian discriminative inference given a model family $p(c, \mathbf{x}, \theta)$ that is assumed to contain all our prior information but still known to be incorrect. This falls in between “standard” Bayesian generative modeling and Bayesian regression, where the margin $p(\mathbf{x}, \theta)$ is known to be uninformative about $p(c|\mathbf{x}, \theta)$. We give an axiomatic proof that *discriminative posterior* is consistent for conditional inference; using the discriminative posterior is standard practice in classical Bayesian regression, but we show that it is theoretically justified for model families of joint densities as well. A practical benefit compared to Bayesian regression is that the standard methods of handling missing values in generative modeling can be extended into discriminative inference, which is useful if the amount of data is small. Compared to standard generative modeling, discriminative posterior results in better conditional inference if the model family is incorrect. If the model family contains also the true model, the discriminative posterior gives the same result as standard Bayesian generative modeling. Practical computation is done with Markov chain Monte Carlo.

1 Introduction

Our aim is Bayesian discriminative inference in the case where the model family $p(c, \mathbf{x}, \theta)$ is known to be incorrect. Here \mathbf{x} is a data vector and c its class, and the θ are parameters of the model family. By discriminative we mean predicting the conditional distribution $p(c | \mathbf{x})$.

The Bayesian approach of using the posterior of the generative model family $p(c, \mathbf{x}, \theta)$ has not been shown to be justified in this case, and it is known that it does not always generalize well to new data (in case of point estimates, see for example

[1, 2, 3]; in this paper we provide a toy example that illustrates the fact for posterior distributions). Therefore alternative approaches such as Bayesian regression are applied [4]. It can be argued that the best solution is to improve the model family by incorporating more prior knowledge. This is not always possible or feasible, however, and simplified models are being generally used, often with good results. For example, it is often practical to use mixture models even if it is known a priori that the data cannot be faithfully described by them (see for example [5]). There are good reasons for still applying Bayesian-style techniques [6] but the general problem of how to best do inference with incorrect model families is still open.

In practice, the usual method for discriminative tasks is Bayesian regression. It disregards all assumptions about the distribution of \mathbf{x} , and considers \mathbf{x} only as covariates of the model for c . Bayesian regression may give superior results in discriminative inference, but the omission of a generative model for \mathbf{x} (although it may be readily available) makes it difficult to handle missing values in the data. Numerous heuristic methods for imputing missing values have been suggested, see for example [7], but no theoretical arguments of their optimality have been presented. Here we assume that we are given a generative model family of the full data (\mathbf{x}, c) , and therefore have a generative mechanism readily available for imputing missing values.

From the generative modeling perspective, Bayesian regression ignores any information about c supplied by the marginal distribution of \mathbf{x} . This is justified if (i) the covariates are explicitly chosen when designing the experimental setting and hence are not noisy, or (ii) there is a separate set of parameters for generating \mathbf{x} on the one hand and c given \mathbf{x} on the other, and the sets are assumed to be independent in their prior distribution. In the latter case the posterior factors out into two parts, and the parameters used for generating \mathbf{x} are neither needed nor useful in the regression task. See for instance [4, 8] for more details. However, there has been no theoretical justification for Bayesian regression in the more general setting where the independence does not hold.

For point estimates of generative models it is well known that maximizing the joint likelihood and the conditional likelihood give in general different results. Maximum conditional likelihood gives asymptotically a better estimate of the conditional likelihood [2], and it can be optimized with expectation-maximization-type procedures [9, 10]. In this paper we extend that line of work to show that the two different approaches, joint and conditional modeling, result in different posterior distributions which are asymptotically equal only if the true model is within the model family. We give an axiomatic justification to the discriminative posterior, and demonstrate empirically that it works as expected. If there are no covariates, the discriminative posterior is the same as the standard posterior of joint density modeling, that is, ordinary Bayesian inference.

To our knowledge the extension from point estimates to a posterior distribution is new. We are aware of only one suggestion, the so-called supervised posterior [11], which also has empirical support in the sense of maximum a posteriori estimates [12]. The posterior has, however, only been justified heuristically.

For the purpose of regression, the discriminative posterior makes it possible to use more general model structures than standard Bayesian regression; in essence any generative model family can be used. In addition to giving a general justification

to Bayesian regression-type modeling, predictions given a generative model family $p(c, \mathbf{x}, \theta)$ should be better if the whole model is (at least approximately) correct. The additional benefit is that the use of the full generative model gives a principled way of handling missing values. The gained advantage, compared to using the standard non-discriminative posterior, is that the predictions should be more accurate assuming the model family is incorrect.

In this paper, we present the necessary background and definitions in Section 2. The discriminative posterior is derived briefly from a set of five axioms in Section 3; the full proof is included as an appendix. There is a close resemblance to Cox axioms, and standard Bayesian inference can indeed be derived also from this set. However, the new axioms allow also inference in the case where the model manifold is known to be inadequate. In section 4 we show that discriminative posterior can be extended in a standard manner [7] to handle data missing at random. In Section 5 we present some experimental evidence that the discriminative posterior behaves as expected.

2 Aims and Definitions

In this paper, we prove the following two claims; the claims follow from Theorem 3.1, discriminative posterior, which is the main result of this paper.

Well-known Given a discriminative model, a model $p(c \mid \mathbf{x}; \theta)$ for the conditional density, Bayesian regression results in consistent conditional inference.

New Given a joint density model $p(c, \mathbf{x} \mid \theta)$, discriminative posterior results in consistent conditional inference.

In accordance with [13], we call inference consistent if the utility is maximized with large data sets. This paper proves both of the above claims. Notice that although the claim 1 is well known, it has not been proven, aside from the special case where the priors for the margin \mathbf{x} and $c \mid \mathbf{x}$ are independent, as discussed in the introduction and in [4].

2.1 Setup

Throughout the paper, observations are denoted by (c, \mathbf{x}) , and assumed to be i.i.d. We use Θ to denote the set of all possible models that could generate the observations. Models that are applicable in practice are restricted to a lower dimensional manifold $\bar{\Theta}$ of models, $\bar{\Theta} \subseteq \Theta$. In other words, the subspace $\bar{\Theta}$ defines our *model family*, in this work denoted by a distribution $p(c, \mathbf{x}, \theta)$ parameterized by $\theta \in \bar{\Theta}$.

There exists a model in Θ which describes the “true” model, which has actually generated the observations and is typically unknown. With slight abuse of notation we denote this model by $\tilde{\theta} \in \Theta$, with the understanding that it may be outside our parametric model family. In fact, in practice no probabilistic model is perfectly true and is false to some extent, that is, the data has usually not been generated by a model in our model family, $\tilde{\theta} \notin \bar{\Theta}$.

The distribution induced in the model parameter space $\overline{\Theta}$ after observing the data $D = \{(c, \mathbf{x})\}_{i=1}^n$ is referred to as a *posterior*. By standard posterior we mean the posterior obtained from Bayes formula using a full joint density model. In this paper we discuss the discriminative posterior, which is obtained from axioms 1–5 below.

2.2 Utility Function and Point Estimates

In this subsection, we introduce the research problem by recapitulating the known difference between point estimates of joint and conditional likelihood. We present the point estimates in terms of Kullback-Leibler divergences, in a form that allows generalizing from point estimates to the posterior distribution in section 3.

Bayesian inference can be derived in a decision theoretic framework as maximization of the expected utility of the decision maker [14]. In general, the choice of the utility function is subjective. However, several arguments for using log-probability as utility function can be made, see for example [14].

When inspecting a full generative model at the limit where the amount of data is infinite, the joint posterior distribution $p_j(\theta | D) \propto p(\theta) \prod_{(c, \mathbf{x}) \in D} p(c, \mathbf{x} | \theta)$ becomes a point solution, $p_j(\theta | D) = \delta(\theta - \hat{\theta})$.¹ An accurate approximation of the log-likelihood is produced by a utility function minimizing the approximation error K_{JOINT} between the point estimate $\hat{\theta}_{JOINT}$ and the true model $\tilde{\theta}$ as follows:

$$\begin{aligned} \hat{\theta}_{JOINT} &= \arg \min_{\theta \in \overline{\Theta}} K_{JOINT}(\tilde{\theta}, \theta) \quad \text{where} \\ K_{JOINT}(\tilde{\theta}, \theta) &= \sum_c \int p(c, \mathbf{x} | \tilde{\theta}) \log \frac{p(c, \mathbf{x} | \tilde{\theta})}{p(c, \mathbf{x} | \theta)} d\mathbf{x}, \end{aligned} \quad (1)$$

If the true model is in the model family, that is, $\tilde{\theta} \in \overline{\Theta}$, equation (1) can be minimized to zero and the resulting point estimate is effectively the MAP solution. If $\tilde{\theta} \notin \overline{\Theta}$ the resulting point estimate is the best estimate of the true joint distribution $p(c, \mathbf{x} | \tilde{\theta})$ with respect to K_{JOINT} .

However, the joint estimate may not be optimal if we are interested in approximating some other quantity than the likelihood. Consider the problem of finding the best point estimate $\hat{\theta}_{COND}$ for the conditional distribution $p(c | \mathbf{x}, \theta)$. The average KL-divergence between the true conditional distribution at $\tilde{\theta}$ and its estimate at θ is given by

$$K_{COND}(\tilde{\theta}, \theta) = \int p(\mathbf{x} | \tilde{\theta}) \sum_c p(c | \mathbf{x}, \tilde{\theta}) \log \frac{p(c | \mathbf{x}, \tilde{\theta})}{p(c | \mathbf{x}, \theta)} d\mathbf{x}, \quad (2)$$

and the best point estimate with respect to K_{COND} is

$$\hat{\theta}_{COND} = \arg \min_{\theta \in \overline{\Theta}} K_{COND}(\tilde{\theta}, \theta). \quad (3)$$

By equations (1) and (2) we may write

$$K_{JOINT}(\tilde{\theta}, \theta) = K_{COND}(\tilde{\theta}, \theta) + \int p(\mathbf{x} | \tilde{\theta}) \log \frac{p(\mathbf{x} | \tilde{\theta})}{p(\mathbf{x} | \theta)} d\mathbf{x}. \quad (4)$$

¹Strictly speaking, the posterior can also have multiple modes; we will not treat these special cases here but they do not restrict the generality.

Therefore the point estimates $\hat{\theta}_{JOINT}$ and $\hat{\theta}_{COND}$ are different in general. If the model that has generated the data does not belong to the model family, that is $\tilde{\theta} \notin \bar{\Theta}$, then by Equation (4) the joint estimate is generally worse than the conditional estimate in conditional inference, measured in terms of conditional likelihood. See also [2].

2.3 Discriminative versus generative models

A discriminative model does not make any assumptions on the distribution of the margin of \mathbf{x} . That is, it does not incorporate a generative model for \mathbf{x} , and can be interpreted to rather use the empirical distribution of $p(\mathbf{x})$ as its margin [15].

A generative model, on the other hand, assumes a specific parametric form for the full distribution $p(c, \mathbf{x}, \theta)$. A generative model can be learned either as a joint density model or in a discriminative manner. Our point in this paper is that the selection corresponds to choosing the utility function; this is actually our fifth axiom in section 3. In joint density modeling the utility is to model the full distribution $p(c, \mathbf{x})$ as accurately as possible, which corresponds to computing the standard posterior incorporating the likelihood function. In discriminative learning the utility is to model the conditional distribution of $p(c | \mathbf{x})$, and the result is a discriminative posterior incorporating the conditional likelihood. A generative model optimized in discriminative manner is referred to as a discriminative joint density model in the following.

For generative models and in case of point estimates, if the model family is correct, a maximum likelihood (ML) solution is better for predicting $p(c|\mathbf{x})$ than maximum conditional likelihood (CML). They have the same maximum, but the asymptotic variance of CML estimate is higher [16]. However, in case of incorrect models, a maximum conditional likelihood estimate is better than maximum likelihood [2]. For an example illustrating that CML can be better than ML in predicting $p(c|\mathbf{x})$, see [17].

Since the discriminative joint density model has a more restricted model structure than Bayesian regression, we expect it to perform better with small amounts of data. More formally, later in Theorem 3.1 we move from point estimate of Equation (3) to a discriminative posterior distribution $p_d(\theta | D)$ over the model parameters θ in $\bar{\Theta}$. Since the posterior is normalized to unity, $\int_{\bar{\Theta}} p_d(\theta | D) d\theta = 1$, the values of the posterior $p_d(\theta | D)$ are generally smaller for larger model families; the posterior is more diffuse. Equation (2) can be generalized to the expectation of the approximation error,

$$E_{p_d(\theta|D)} \left[K_{COND}(\tilde{\theta}, \theta) \right] = - \int \sum_c p(\mathbf{x}, c | \tilde{\theta}) p_d(\theta | D) \log p(c | \mathbf{x}, \theta) d\mathbf{x} d\theta + \text{const.} \quad (5)$$

The expected approximation error is small when both the discriminative posterior distribution $p_d(\theta | D)$ and the conditional likelihood $p(c | \mathbf{x}, \theta)$ are large at the same time. For small amounts of data, if the model family is too large, the values of the posterior $p_d(\theta | D)$ are small. The discriminative joint density model has a more restricted model family than that of the Bayesian regression, and hence the values of the posterior are larger. If the model is approximately correct, the discriminative joint density model will have a smaller approximation error than the Bayesian regression, that is, $p(c | \mathbf{x}, \theta)$ is large somewhere in $\bar{\Theta}$. This is analogous to selecting the model family that maximizes the evidence (in our case the expected conditional log-likelihood) in

Bayesian inference; choosing a model family that is too complex leads to small evidence (see, e.g., [18]). The difference to the traditional Bayesian inference is that we do not require that the true data generating distribution $\hat{\theta}$ is contained in the parameter space $\overline{\Theta}$ under consideration.

3 Axiomatic Derivation of Discriminative Posterior

In this section we generalize the point estimate $\hat{\theta}_{COND}$ presented in section 2.2 to a *discriminative posterior distribution* over $\theta \in \overline{\Theta}$.

Theorem 3.1 (Discriminative posterior distribution) *It follows from axioms 1–6 listed below that, given data $D = \{(c_i, \mathbf{x}_i)\}_{i=1}^n$, the discriminative posterior distribution $p_d(\theta | D)$ is of the form*

$$p_d(\theta | D) \propto p(\theta) \prod_{(c, \mathbf{x}) \in D} p(c | \mathbf{x}, \theta) .$$

The predictive distribution for new $\tilde{\mathbf{x}}$, obtained by integrating over this posterior, $p(c | \tilde{\mathbf{x}}, D) = \int p_d(\theta | D) p(c | \tilde{\mathbf{x}}, \theta) d\theta$, is consistent for conditional inference. That is, p_d is consistent for the utility of conditional likelihood.

The discriminative posterior follows from requiring the following axioms to hold:

1. The posterior $p_d(\theta | D)$ can be represented by non-negative real numbers that satisfy $\int_{\overline{\Theta}} p_d(\theta | D) d\theta = 1$.
2. A model $\theta \in \overline{\Theta}$ can be represented as a function $h((c, \mathbf{x}), \theta)$ that maps the observations (c, \mathbf{x}) to real numbers.
3. The posterior, after observing a data set D followed by an observation (c, \mathbf{x}) , is given by $p_d(\theta | D \cup (c, \mathbf{x})) = F(h((c, \mathbf{x}), \theta), p_d(\theta | D))$, where F is a twice differentiable function in both of its parameters.
4. Exchangeability: The value of the posterior is independent of the ordering of the observations. That is, the posterior after two observations (c, \mathbf{x}) and (c', \mathbf{x}') is the same irrespective of their ordering:
 $F(h((c', \mathbf{x}'), \theta), p_d(\theta | (c, \mathbf{x}) \cup D)) = F(h((c, \mathbf{x}), \theta), p_d(\theta | (c', \mathbf{x}') \cup D))$.
5. The posterior must agree with the utility. For $\tilde{\theta} \in \Theta$, and $\theta_1, \theta_2 \in \overline{\Theta}$, the following condition is satisfied:

$$p_d(\theta_1 | D_{\tilde{\theta}}) \leq p_d(\theta_2 | D_{\tilde{\theta}}) \Leftrightarrow K_{COND}(\tilde{\theta}, \theta_1) \geq K_{COND}(\tilde{\theta}, \theta_2) ,$$

where $D_{\tilde{\theta}}$ is a very large data set sampled from $p(c, \mathbf{x} | \tilde{\theta})$. We further assume that the discriminative posteriors p_d at θ_1 and θ_2 are equal only if the corresponding conditional KL-divergences K_{COND} are equal.

The first axiom above is simply a requirement that the posterior is a probability distribution in the parameter space.

The second axiom defines a model in general terms; we define it as a mapping from event space into a real number.

The third axiom makes smoothness assumptions on the posterior. The reason for the axiom is technical; the smoothness is used in the proofs. Cox [19] makes similar assumptions, and our proof therefore holds in the same scope as Cox’s; see [20].

The fourth axiom requires exchangeability; the shape of the posterior distribution should not depend on the order in which the observations are made. This deviates slightly from analogous earlier proofs for standard Bayesian inference, which have rather used the requirement of associativity [19] or included it as an additional constraint in modeling after presenting the axioms [14].

The fifth axiom states, in essence, that asymptotically (at the limit of a large but finite data set) the shape of the posterior is such that the posterior is always smaller if the “distance” $K_{COND}(\tilde{\theta}, \theta)$ is larger. If the opposite would be true, the integral over the discriminative posterior would give larger weight to solutions further away from the true model, leading to a larger error measured by $K_{COND}(\tilde{\theta}, \theta)$.

Axioms 1–5 are sufficient to fix the discriminative posterior, up to a monotonic transformation. To fix the monotonic transformation we introduce the sixth axiom:

6. For fixed \mathbf{x} the model reduces to the standard posterior. For the data set $D_x = \{(c, \mathbf{x}') \in D \mid \mathbf{x}' = \mathbf{x}\}$, the discriminative posterior $p_d(\theta \mid D_x)$ matches the standard posterior

$$p^x(c \mid \theta) \equiv p(c \mid \mathbf{x}, \theta).$$

We use $p(\theta) = p_d(\theta \mid \emptyset)$ to denote the posterior when no data is observed (*prior distribution*).

Proof The proof is in the Appendix. For clarity of presentation, we additionally sketch the proof in the following.

Proposition 3.2 (*F is isomorphic to multiplication*) *It follows from axiom 4 that the function F is of the form*

$$f(F(h((c, \mathbf{x}), \theta), p_d(\theta \mid D))) \propto h((c, \mathbf{x}), \theta) f(p_d(\theta \mid D)), \quad (6)$$

where f is a monotonic function which we, by convention, fix to the identity function.²

The problem then reduces to finding a functional form for $h((c, \mathbf{x}), \theta)$. Utilizing both the equality part and the inequality part of axiom 5, the following proposition can be derived.

Proposition 3.3 *It follows from axiom 5 that*

$$h((c, \mathbf{x}), \theta) \propto p(c \mid \mathbf{x}, \theta)^A \quad \text{where } A > 0. \quad (7)$$

Finally, the axiom 6 effectively states that we decide to follow the Bayesian convention for a fixed \mathbf{x} , that is, to set $A = 1$.

²We follow here Cox [19] and subsequent work, see e.g. [20]. A difference which does not affect the need for the convention is that usually multiplicativity is derived based on the assumption of associativity, not exchangeability. Our setup is slightly different from Cox, since instead of updating beliefs on events, we here consider updating beliefs in a family of models.

4 Modeling Missing Data

New Discriminative posterior gives a theoretically justified way of handling missing values in discriminative tasks.

Discriminative models cannot readily handle components missing from the data vector, since the data is used only as covariates. However, standard methods of handling missing data with generative models [7] can be applied with the discriminative posterior.

The additional assumption we need to make is a model for which data are missing. Below we derive the formulas for the common case of data missing independently at random. Extensions incorporating prior information of the process by which the data are missing are straightforward, although possibly not trivial.

Write the observations $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2)$. Assume that \mathbf{x}_1 can be missing and denote a missing observation by $\mathbf{x}_1 = \emptyset$. The task is still to predict c by $p(c \mid \mathbf{x}_1, \mathbf{x}_2)$.

Since we are given a model for the joint distribution which is assumed to be approximately correct, it will be used to model the missing data. We denote this by $q(c, \mathbf{x}_1, \mathbf{x}_2 \mid \theta')$, $\mathbf{x}_1 \neq \emptyset$, with a prior $q(\theta')$, $\theta' \in \overline{\Theta}'$. We further denote the parameters of the missing data mechanism by λ . Now, similar to joint density modeling, if the priors for the joint model, $q(\theta')$, and missing-data mechanism, $g(\lambda)$, are independent, the missing data mechanism is ignorable [7]. In other words, the posterior that takes the missing data into account can be written as

$$p_d(\theta \mid D) \propto q(\theta')g(\lambda) \prod_{\mathbf{y} \in D_{full}} q(c \mid \mathbf{x}_1, \mathbf{x}_2, \theta') \prod_{\mathbf{y} \in D_{missing}} q(c \mid \mathbf{x}_2, \theta'), \quad (8)$$

where $\mathbf{y} = (c, \mathbf{x}_1, \mathbf{x}_2)$ and we have used D_{full} and $D_{missing}$ to denote the portions of data set with $\mathbf{x}_1 \neq \emptyset$ and $\mathbf{x}_1 = \emptyset$, respectively.

Equation (8) has been obtained by using q to construct a model family in which the data is missing independently at random with probability λ , having a prior $g(\lambda)$. That is, we define a model family that generates the missing data in addition to the non-missing data,

$$p(c, \mathbf{x}_1, \mathbf{x}_2 \mid \theta) \equiv \begin{cases} (1 - \lambda)q(c, \mathbf{x}_1, \mathbf{x}_2 \mid \theta') & , \quad \mathbf{x}_1 \neq \emptyset \\ \lambda \int_{\text{supp}(\mathbf{x}_1) \setminus \emptyset} q(c, \mathbf{x}, \mathbf{x}_2 \mid \theta') d\mathbf{x} = \lambda q(c, \mathbf{x}_2 \mid \theta') & , \quad \mathbf{x}_1 = \emptyset, \end{cases} \quad (9)$$

where $\theta \in \overline{\Theta}$ and $\text{supp}(\mathbf{x}_1)$ denotes the support of \mathbf{x}_1 . The parameter space $\overline{\Theta}$ is spanned by $\theta' \in \overline{\Theta}'$ and λ . The equation (8) follows directly by applying Theorem 3.1 to $p(c \mid \mathbf{x}_1, \mathbf{x}_2, \theta)$.³ Notice that the posterior for θ' of (8) is independent of λ . The division to \mathbf{x}_1 and \mathbf{x}_2 can be made separately for each data item. The items can have different numbers of missing components, each component k having a probability λ_k of being missing.

³Notice that $p(c \mid \mathbf{x}_1, \mathbf{x}_2, \theta) = q(c \mid \mathbf{x}_1, \mathbf{x}_2, \theta')$ when $\mathbf{x}_1 \neq \emptyset$, and $p(c \mid \mathbf{x}_1 = \emptyset, \mathbf{x}_2, \theta) = q(c \mid \mathbf{x}_2, \theta')$.

5 Experiments

5.1 Implementation of Discriminative Sampling

The discriminative posterior can be sampled with an ordinary Metropolis-Hastings algorithm where the standard posterior $p(\theta \mid D) \propto \prod_{i=1}^n p(c_i, x_i \mid \theta)p(\theta)$ is simply replaced by the discriminative version $p_d(\theta \mid D) \propto \prod_{i=1}^n p(c_i \mid x_i, \theta)p(\theta)$, where $p(c_i \mid x_i, \theta) = \frac{p(c_i, x_i \mid \theta)}{p(x_i \mid \theta)}$.

In MCMC sampling of the discriminative posterior, the normalization term of the conditional likelihood poses problems, since it involves a marginalization over the class variable c and latent variables z , that is, $p(\mathbf{x}_i \mid \theta) = \sum_c \int_{\text{supp}(z)} p(\mathbf{x}_i, z, c \mid \theta) dz$. In case of discrete latent variables, such as in mixture models, the marginalization reduces to simple summations and can be computed exactly and efficiently. If the model contains continuous latent variables the integral needs to be evaluated numerically.

5.2 Performance as a function of the incorrectness of the model

We first set up a toy example where the distance between the model family and the true model is varied.

5.2.1 Background

In this experiment we compare the performance of logistic regression and a mixture model. Logistic regression was chosen since many of the discriminative models can be seen as extensions of it, for example conditional random fields [8]. In case of logistic regression, the following theorem exists:

Theorem 5.1 [21]

For a k -class classification problem with equal priors, the pairwise log-odds ratio of the class posteriors is affine if and only if the class conditional distributions belong to any fixed exponential family.

That is, the model family of logistic regression incorporates the model family of any generative exponential (mixture) model⁴. A direct interpretation is that any generative exponential family mixture model defines a smaller subspace within the parametric space of the logistic regression model.

Since the model family of logistic regression is larger, it is asymptotically better, but if the number of data is small, generative models can be better due to the restricted parameter space; see for example [3] for a comparison of naive Bayes and logistic regression. This happens if the particular model family is at least approximately correct.

5.2.2 Experiment

The true model generates ten-dimensional data \mathbf{x} from the two-component mixture $p(\mathbf{x}) = \sum_{j=1}^2 \frac{1}{2} \times p(\mathbf{x} \mid \tilde{\mu}_j, \tilde{\sigma}_j)$, where j indexes the mixture component, and $p(\mathbf{x} \mid$

⁴Banerjee [21] provides a proof also for conditional random fields.

$\tilde{\mu}_j, \tilde{\sigma}_j$) is a Gaussian with mean $\tilde{\mu}_j$ and standard deviation $\tilde{\sigma}_j$. The data is labeled according to the generating mixture component (i.e., the “class” variable) $j \in \{1, 2\}$. Two of the ten dimensions contain information about the class. The “true” parameters, used to generate the data, on these dimensions are $\tilde{\mu}_1 = 5$, $\tilde{\sigma}_1 = 2$, and $\tilde{\mu}_2 = 9$, $\tilde{\sigma}_2 = 2$. The remaining eight dimensions are Gaussian noise with $\tilde{\mu} = 9$, $\tilde{\sigma} = 2$ for both components.

The “incorrect” generative model used for inference is a mixture of two Gaussians where the variances are constrained to be $\sigma_j = k \cdot \mu_j + 2$. With increasing k the model family thus draws further away from the true model. We assume for both of the μ_j a Gaussian prior $p(\mu_j | m, s)$ having the same fixed hyperparameters⁵ $m = 7$, $s = 7$ for all dimensions.

Data sets were generated from the true model with 10000 test data points and a varying number of training data points ($N_D = \{32, 64, 128, 256, 512, 1024\}$). Both the discriminative and standard posteriors were sampled for the model. The goodness of the models was evaluated by perplexity of conditional likelihood on the test data set.

Standard and discriminative posteriors of the incorrect generative model are compared against Bayesian logistic regression. A uniform prior for the parameters β of the logistic regression model was assumed. As mentioned in subsection 5.2.1 above, the Gaussian naive Bayes and logistic regression are connected (see for example [22] for exact mapping between the parameters). The parameter space of the logistic model incorporates the true model as well as the incorrect model family, and is therefore the optimal discriminative model for this toy data. However, since the parameter space of the logistic regression model is larger, the model is expected to need more data samples for good predictions.

The models perform as expected (Figure 1). The model family of our incorrect model was chosen such that it contains useful prior knowledge about the distribution of \mathbf{x} . Compared with logistic regression, the model family becomes more restricted which is beneficial for small learning data sets (see Figure 1a). Compared to joint density sampling, discriminative sampling results in better predictions of the class c when the model family is incorrect.

The models were compared more quantitatively by repeating the sampling ten times for a fixed value of $k = 2$ for each of the learning data set sizes; in every repeat also a new data set was generated. The results in Figure 2 confirm the qualitative findings of Figure 1.

The posterior was sampled with Metropolis-Hastings algorithm with a Gaussian jump kernel. In case of joint density and discriminative sampling, three chains were sampled with a burn-in of 500 iterations each, after which every fifth sample was collected. Bayesian regression required more samples for convergence, so a burn-in of 5500 samples was used. Convergence was estimated when carrying out experiments for Figure 2; the length of sampling chains was set such that the confidence intervals for each of the models was roughly the same. The total number of samples was 900 per chain. The width of jump kernel was chosen as a linear function of data such that the acceptance rate was between 0.2–0.4 [4] as the amount of data was increased. Selection was carried out by preliminary tests with different random seeds.

⁵The model is thus slightly incorrect even with $k = 0$.

Missing Data. The experiment with toy data is continued in a setting where 50% of the learning data are missing at random.

MCMC sampling was carried out as described above. For the logistic regression model, a multiple imputation scheme is applied, as recommended in [4]. In order to have sampling conditions comparable to sampling from a generative model, each sampling chain used one imputed data set. Imputation was carried out with the generative model (representing our current best knowledge, differing from the true joint model, however).

As can be seen from Figure 1, discriminative sampling is better than joint density sampling when the model is incorrect. The performance of Bayesian regression seems to be affected heavily by the incorrect generative model used to generate missing data. As can be seen from Figure 2, surprisingly the Bayesian regression is even worse than standard posterior. The performance could be increased by imputing more than one data sets with the cost of additional computational complexity, however.

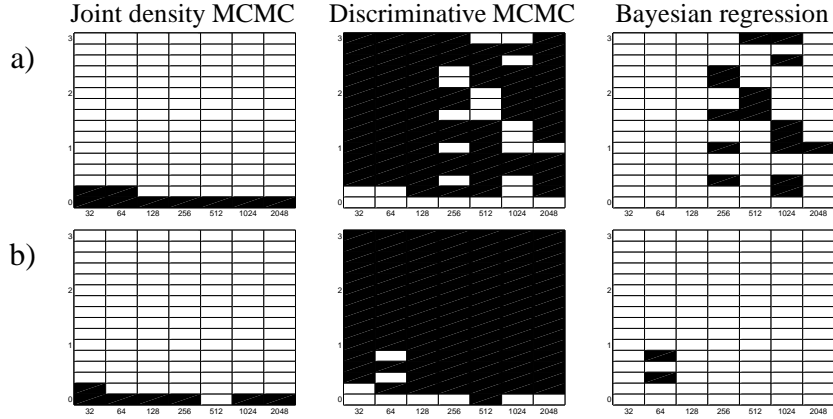


Figure 1: A comparison of joint density sampling, discriminative sampling, and Bayesian regression, a) Full data, b) 50% of the data missing. Grid points where the method is better than others are marked with black. The X-axis denotes the amount of data, and Y-axis the deviance of the the model family from the “true” model (i.e., the value of k). The methods are prone to sampling error, but the following general conclusions can be made: Bayesian generative modeling (“Joint density MCMC”) is best when the model family is approximately correct. Discriminative posterior (“Discriminative MCMC”) is better when the model is incorrect and the learning data set is small. As the amount of data is increased, Bayesian regression and discriminative posterior show roughly equal performance (see also Figure 2).

5.3 Document Modeling

As a demonstration in a practical domain we applied the discriminative posterior to document classification. We used the Reuters data set [23], of which we selected a

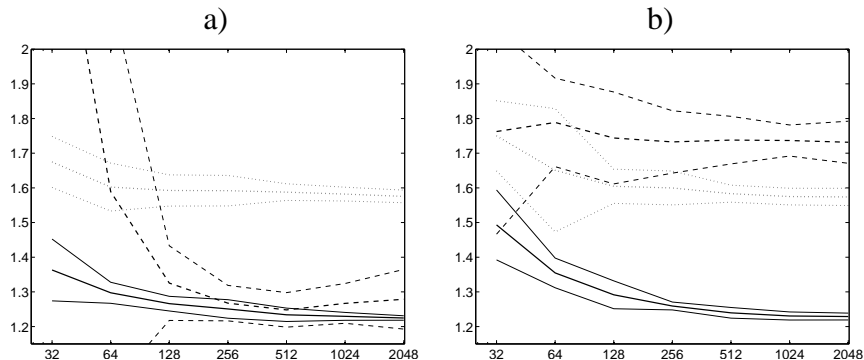


Figure 2: A comparison of joint density sampling (dotted line), discriminative sampling (solid line), and Bayesian regression (dashed line) with an incorrect model. X-axis: learning data set size, Y-axis: perplexity. Also the 95 % confidence intervals are plotted (with the thin lines). Logistic regression performs significantly worse than discriminative posterior with small data set sizes, whereas with large data sets the performance is roughly equal. Joint density modeling is consistently worse. The model was fixed to $k = 2$, ten individual runs were carried out in order to compute 95% confidence intervals. a) Full data, b) 50% of the data missing.

subset of 1100 documents from four categories (CCAT, ECAT, GCAT, MCAT). Each selected document was classified to exactly one of the four classes. As a preprocessing stage, we chose the 25 most informative words within the training set of 100 documents, having the highest mutual information between classes and words [24]. The remaining 1000 documents were used as a test set.

We first applied a mixture of unigrams model (MUM, see Figure 3 left) [25]. The model assumes that each document is generated from a mixture of M hidden “topics,” $p(\mathbf{x}_i | \theta) = \sum_{j=1}^M \pi(j) p(\mathbf{x}_i | \beta_j)$, where j is the index of the topic, and β_j the multinomial parameters that generate words from the topic. The vector \mathbf{x}_i contains the observed word counts (with a total of N_W) for document i , and $\pi(j)$ is the probability of generating words from the topic j . The usual approach [4] for modeling paired data $\{\mathbf{x}_i, c_i\}_{i=1}^{N_D}$ by a joint density mixture model was applied; c was associated with the label of the mixture component from which the data is assumed to be generated. Dirichlet priors were used for the β, π , with hyperparameters set to 25 and 1, respectively. We used the simplest form of MUM containing one topic vector per class. The sampler used Metropolis-Hastings with a Gaussian jump kernel. The kernel width was chosen such that the acceptance rate was roughly 0.2 [4].

As an example of a model with continuous hidden variables, we implemented the Latent Dirichlet Allocation (LDA) or discrete PCA model [26, 27]. We constructed a variant of the model that generates also the classes, shown in Figure 3 (right). The topology of the model is a mixture of LDAs; the generating mixture component z_c is first sampled from π_c . The component indexes a row in the matrix α , and (for simplicity) contains a direct mapping to c . Now, given α , the generative model for

words is an ordinary LDA. The π is a topic distribution drawn individually for each document d from a Dirichlet with parameters $\alpha(z_c)$, that is, $\text{Dirichlet}(\alpha(z_c))$. Each word w_{nd} belongs to one topic z_{nd} which is picked from $\text{Multinomial}(\pi)$. The word is then generated from $\text{Multinomial}(\beta(z_{nd}, \cdot))$, where z_{nd} indexes a row in the matrix β . We assume a Dirichlet prior for the β , with hyperparameters set to 2, Dirichlet prior for the α , with hyperparameters set to 1, and a Dirichlet prior for π_c with hyperparameters equal to 50. The parameter values were set in initial test runs (with a separate data set from the Reuters corpus). Four topic vectors were assumed, making the model structure similar to [28]. The difference to MUM is that in LDA-type models a document can be generated from several topics.

Sampling was carried out using Metropolis-Hastings with a Gaussian jump kernel, where the kernel width was chosen such that the acceptance rate was roughly 0.2 [4]. The necessary integrals were computed with Monte Carlo integration. The convergence of integration was monitored with a jackknife estimate of standard error [29]; sampling was ended when the estimate was less than 5 % of the value of the integral. The length of burn-in was 100 iterations, after which every tenth sample was picked. The total number of collected samples was 100. The probabilities were clipped to the range $[e^{-22}, 1]$.

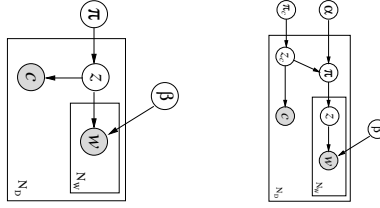


Figure 3: Left: Mixture of Unigrams. Right: Mixture of Latent Dirichlet Allocation models.

5.3.1 Results

Discriminative sampling is better for both models (Table 1).

Table 1: Comparison of sampling from the ordinary posterior (jMCMC) and discriminative posterior (dMCMC) for two model families: Mixture of Unigrams (MUM) and mixture of Latent Dirichlet Allocation models (mLDA). The figures are perplexities of the 1000 document test set; there were 100 learning data points. Small perplexity is better; random guessing gives perplexity of 4.

Model	dMCMC	jMCMC	Conditional ML
MUM	2.56	3.98	4.84
mLDA	2.36	3.92	3.14

6 Discussion

We have introduced a principled way of making conditional inference with a discriminative posterior. Compared to standard joint posterior density, discriminative posterior results in better inference when the model family is incorrect, which is usually the case. Compared to purely discriminative modeling, discriminative posterior is better in case of small data sets if the model family is at least approximately correct. Additionally, we have introduced a justified method for incorporating missing data into discriminative modeling.

Joint density modeling, discriminative joint density modeling, and Bayesian regression can be seen as making different assumptions on the margin model $p(\mathbf{x} \mid \theta)$. Joint density modeling assumes the model family to be correct, and hence also the model of \mathbf{x} margin to be correct. If this assumption holds, the discriminative posterior and joint density modeling will asymptotically give the same result. On the other hand, if the assumption does not hold, the discriminative joint density modeling will asymptotically give better or at least as good results. Discriminative joint density modeling assumes that the margin model $p(\mathbf{x} \mid \theta)$ may be incorrect, but the conditional model $p(c \mid \mathbf{x}, \theta)$, derived from the joint model that includes the model for the margin, is in itself at least approximately correct. Then inference is best made with discriminative posterior as in this paper. Finally, if the model family is completely incorrect — or if there is lots of data — a larger, discriminative model family and Bayesian regression should be used.

Another approach to the same problem was suggested in [30], where the traditional generative view to discriminative modeling has been extended by complementing the conditional model for $p(c \mid \mathbf{x}, \theta)$ with a model for $p(\mathbf{x} \mid \theta')$, to form the joint density model $p(c, \mathbf{x} \mid \theta, \theta') = p(c \mid \mathbf{x}, \theta)p(\mathbf{x} \mid \theta')$. That is, a larger model family is postulated with additional parameters θ' for modeling the marginal \mathbf{x} . The conditional density model $p(c \mid \mathbf{x}, \theta)$ is derived by Bayes rule from a formula for the joint density, $p(c, \mathbf{x} \mid \theta)$, and the model for the marginal $p(\mathbf{x} \mid \theta')$ is obtained by marginalizing it.

This conceptualization is very useful for semisupervised learning; the dependency between θ and θ' can be tuned by choosing a suitable prior, which allows balancing between discriminative and generative modeling. The optimum for semisupervised learning is found in between the two extremes. The approach of [30] contains our discriminative posterior distribution as a special case in the limit where the priors are independent, that is, $p(\theta, \theta') = p(\theta)p(\theta')$ where the parameters θ and θ' can be treated independently.

Also [30] can be viewed as giving a theoretical justification for Bayesian discriminative learning, based on generative modeling. The work introduces a method of extending a fully discriminative model into a generative model, making discriminative learning a special case of optimizing the likelihood (that is, the case where priors separate). Our work starts from different assumptions. We assume that the utility functions can be different, depending on the goal of the modeler. As we show in this paper, the requirement of our axiom 5 that the utility should agree with the posterior will eventually lead us to the proper form of the posterior. Here we chose conditional inference as the utility, obtaining a discriminative posterior. If the utility had been joint modeling of (c, \mathbf{x}) , we would have obtained the standard posterior (the case of no covariates). If the given model is incorrect the different utilities lead to different learning schemes. From

this point of view, the approach of [30] is principled only if the “true” model belongs to the postulated larger model family $p(c, \mathbf{x} \mid \theta, \theta')$.

As a practical matter, efficient implementation of sampling from the discriminative posterior will require more work. The sampling schemes applied in this paper are simple and computationally intensive; there exist several more advanced methods which should be tested.

Acknowledgments

This work was supported in part by the PASCAL2 Network of Excellence of the European Community. This publication only reflects the authors’ views.

References

- [1] B. Efron. The efficiency of logistic regression compared to normal discriminant analysis. *Journal of the American Statistical Association*, 70:892–898, 1975.
- [2] Arthur Nádas, David Nahamoo, and Michael A. Picheny. On a model-robust training method for speech recognition. *IEEE transactions on Acoustics, Speech, and Signal Processing*, 39(9):1432–1436, 1988.
- [3] Andrew Y. Ng and Michael I. Jordan. On discriminative vs. generative classifiers: A comparison of logistic regression and naive Bayes. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, pages 841–848. MIT Press, Cambridge, MA, 2002.
- [4] Andrew Gelman, John B. Carlin, Hal S. Stern, and Donald B. Rubin. *Bayesian Data Analysis (2nd edition)*. Chapman & Hall/CRC, Boca Raton, FL, 2003.
- [5] Geoffrey J. McLachlan and David Peel. *Finite Mixture Models*. John Wiley & Sons, New York, NY, 2000.
- [6] L. K. Hansen. Bayesian averaging is well-tempered. In Sara A. Solla, Todd K. Leen, and Klaus-Robert Müller, editors, *Advances in Neural Information Processing Systems 12*, pages 265–271. MIT Press, 2000.
- [7] Roderick J.A. Little and Donald B. Rubin. *Statistical analysis with missing data*. John Wiley & Sons, Inc., Hoboken, NJ, 2nd edition, 2002.
- [8] Charles Sutton and Andrew McCallum. An introduction to conditional random fields for relational learning. In Lise Getoor and Ben Taskar, editors, *Introduction to Statistical Relational Learning*, chapter 4. MIT Press, 2006.
- [9] Jarkko Salojärvi, Kai Puolamäki, and Samuel Kaski. Expectation maximization algorithms for conditional likelihoods. In Luc De Raedt and Stefan Wrobel, editors, *Proceedings of the 22nd International Conference on Machine Learning (ICML-2005)*, pages 753–760, New York, USA, 2005. ACM press.

- [10] P. C. Woodland and D. Povey. Large scale discriminative training of hidden Markov models for speech recognition. *Computer Speech and Language*, 16:25–47, 2002.
- [11] Peter Grünwald, Petri Kontkanen, Petri Myllymäki, Teemu Roos, Henry Tirri, and Hannes Wettig. Supervised posterior distributions. Presentation at the Seventh Valencia International Meeting on Bayesian Statistics, Tenerife, Spain, 2002. <http://homepages.cwi.nl/~pdg/presentationpage.html>.
- [12] Jesús Cerquides and Ramon López Mántaras. Robust Bayesian linear classifier ensembles. In João Gama, Rui Camacho, Pavel Brazdil, Alípio Jorge, and Luís Torgo, editors, *Machine Learning: ECML 2005*, pages 72–83, Berlin, Germany, 2005. Springer-Verlag.
- [13] V.N. Vapnik. *The Nature of Statistical Learning Theory*. Springer, 1995.
- [14] José M. Bernardo and Adrian F. M. Smith. *Bayesian Theory*. John Wiley & Sons Ltd, Chichester, England, 2000.
- [15] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer, New York, 2001.
- [16] Arthur Nádas. A decision theoretic formulation of a training problem in speech recognition and a comparison of training by unconditional versus conditional maximum likelihood. *IEEE transactions on Acoustics, Speech, and Signal Processing*, 31(4):814–817, 1983.
- [17] Tony Jebara and Alex Pentland. On reversing Jensen’s inequality. In Todd K. Leen, Thomas G. Dietterich, and Volker Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 231–237, Cambridge, MA, April 2001. MIT Press.
- [18] Christopher M. Bishop. *Pattern recognition and machine learning*. Springer, 2006.
- [19] R. T. Cox. Probability, frequency and reasonable expectation. *American Journal of Physics*, 17:1–13, 1946.
- [20] J.Y Halpern. A counter example to theorems of Cox and Fine. *Journal of Artificial Intelligence Research*, 10:67–85, 1999.
- [21] Arindam Banerjee. An analysis of logistic models: Exponential family connections and online performance. In *Proceedings of the 2007 SIAM International Conference on Data Mining*, pages 204–215, 2007.
- [22] Tom Mitchell. *Machine Learning*, chapter Generative and Discriminative Classifiers: Naive Bayes and Logistic Regression. 2006. Additional chapter in a draft for the second edition of the book, available from <http://www.cs.cmu.edu/~tom/NewChapters.html>.

- [23] D. D. Lewis, Y. Yang, T. Rose, and F. Li. Rcv1: A new benchmark collection for text categorization research. *Journal of Machine Learning Research*, 5:361–397, 2004.
- [24] Jaakko Peltonen, Janne Sinkkonen, and Samuel Kaski. Sequential information bottleneck for finite data. In Russ Greiner and Dale Schuurmans, editors, *Proceedings of the Twenty-First International Conference on Machine Learning (ICML 2004)*, pages 647–654. Omnipress, Madison, WI, 2004.
- [25] Kamal Nigam, Andrew K. McCallum, Sebastian Thrun, and Tom M. Mitchell. Text classification from labeled and unlabeled documents using EM. *Machine Learning*, 39(2-3):103–134, 2000.
- [26] D. Blei, A. Y. Ng, and M. I. Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.
- [27] Wray Buntine and Aleks Jakulin. Discrete components analysis. In C. Saunders, M. Grobelnik, S. Gunn, and J. Shawe-Taylor, editors, *Subspace, Latent Structure and Feature Selection Techniques*, pages 1–33. Springer-Verlag, 2006.
- [28] Li Fei-Fei and Pietro Perona. A Bayesian hierarchical model for learning natural scene categories. In *Proceedings 2005 IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 524–531, Los Alamitos, CA, USA, 2005. IEEE Computer Society.
- [29] B. Efron and R. Tibshirani. *An Introduction to the Bootstrap*. Chapman&Hall, New York, 1993.
- [30] Julia A. Lasserre, Christopher M. Bishop, and Thomas P. Minka. Principled hybrids of generative and discriminative models. In *Proceedings 2006 IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, pages 87–94, Los Alamitos, CA, USA, 2006. IEEE Computer Society.

A Proofs

We use the notation $\mathbf{r} = (c, \mathbf{x})$ and denote the set of all possible observations \mathbf{r} by R .

For purposes of some parts of the proof, we assume that the set of possible observations R is finite. This assumption is not excessively restrictive, since any well-behaving infinite set of observations and respective probabilistic models can be approximated with an arbitrary accuracy by discretizing R to sufficiently many bins and converting the probabilistic models to the corresponding multinomial distributions over R .

Example: Assume the observations r are real numbers in compact interval $[a, b]$ and they are modelled by a well-behaving probability density $p(r)$, that is, the set of possible observations R is an infinite set. We can approximate the distribution $p(x)$ by partitioning the interval $[a, b]$ into N bins $I(i) = [a + (i - 1)(b - a)/N, a + i(b - a)/N]$ of width $(b - a)/N$ each, where $i \in \{1, \dots, N\}$, and assigning each bin a multinomial probability $\theta_i = \int_{I(i)} p(r)dr$. One possible choice for the family of models $p(r)$ would be the Gaussian distributions parametrized by the mean and variance; the parameter space of these Gaussian distributions would span a 2-dimensional subspace $\bar{\Theta}$ in the $N - 1$ dimensional parameter space Θ of the multinomial distributions.

A.1 From exchangeability $F(h(c', \mathbf{x}'), p_d(\theta \mid (c, \mathbf{x}))) = F(h(c, \mathbf{x}), p_d(\theta \mid (c', \mathbf{x}')))$ it follows that the posterior is homomorphic to multiplicativity

4. Exchangeability: The value of the posterior is independent of the ordering of the observations. That is, the posterior after two observations (c, \mathbf{x}) and (c', \mathbf{x}') is the same irrespective of their ordering: $F(h((c', \mathbf{x}'), \theta), p_d(\theta \mid (c, \mathbf{x}) \cup D)) = F(h((c, \mathbf{x}), \theta), p_d(\theta \mid (c', \mathbf{x}') \cup D))$.

Proof For simplicity let us denote $x = h(c', \mathbf{x}')$, $y = h(c, \mathbf{x})$, and $z = p(\theta)$. The exchangeability axiom thus reduces to the problem of finding a function F such that

$$F(x, F(y, z)) = F(y, F(x, z)) \quad , \quad (10)$$

where $F(y, z) = p_d(\theta \mid (c, \mathbf{x}))$. By denoting $F(x, z) = u$ and $F(y, z) = v$, the equation becomes $F(x, v) = F(y, u)$.

We begin by assuming that function F is differentiable in both its arguments (in a similar manner to Cox). Differentiating with respect to z , y and x in turn, and writing $F_1(p, q)$ for $\frac{\partial F(p, q)}{\partial p}$ and $F_2(p, q)$ for $\frac{\partial F(p, q)}{\partial q}$, we obtain

$$F_2(x, v) \frac{\partial v}{\partial z} = F_2(y, u) \frac{\partial u}{\partial z} \quad (11)$$

$$F_2(x, v) \frac{\partial v}{\partial y} = F_1(y, u) \quad (12)$$

$$F_1(x, v) = F_2(y, u) \frac{\partial u}{\partial x} \quad . \quad (13)$$

Differentiating equation (11) wrt. z , x and y in turn, we get

$$F_{22}(x, v) \left(\frac{\partial v}{\partial z} \right)^2 + F_2(x, v) \frac{\partial^2 v}{\partial z^2} = F_{22}(y, u) \left(\frac{\partial u}{\partial z} \right)^2 + F_2(y, u) \frac{\partial^2 u}{\partial z^2} \quad (14)$$

$$F_{12}(x, v) \frac{\partial v}{\partial z} = F_{22}(y, u) \frac{\partial u}{\partial x} \frac{\partial u}{\partial z} + F_2(y, u) \frac{\partial^2 u}{\partial x \partial z} \quad (15)$$

$$F_{22}(x, v) \frac{\partial v}{\partial y} \frac{\partial v}{\partial z} + F_2(x, v) \frac{\partial^2 v}{\partial z \partial y} = F_{12}(y, u) \frac{\partial u}{\partial z} \quad , \quad (16)$$

and differentiating equation (12) wrt. x we get

$$F_{12}(x, v) \frac{\partial v}{\partial y} = F_{12}(y, u) \frac{\partial u}{\partial x} \quad . \quad (17)$$

By solving $F_{12}(x, v)$ from equation (15) and $F_{12}(y, u)$ from equation (16) and inserting into equation (17), we get

$$\begin{aligned} \frac{F_{22}(y, u) \frac{\partial u}{\partial x} \frac{\partial u}{\partial z} + F_2(y, u) \frac{\partial^2 u}{\partial x \partial z}}{\frac{\partial v}{\partial z}} \frac{\partial v}{\partial y} &= \frac{F_{22}(x, v) \frac{\partial v}{\partial y} \frac{\partial v}{\partial z} + F_2(x, v) \frac{\partial^2 v}{\partial z \partial y}}{\frac{\partial u}{\partial z}} \frac{\partial u}{\partial x} \\ &\Leftrightarrow \\ F_{22}(y, u) \left(\frac{\partial u}{\partial z} \right)^2 - F_{22}(x, v) \left(\frac{\partial v}{\partial z} \right)^2 &= \frac{F_2(x, v) \frac{\partial^2 v}{\partial z \partial y} \frac{\partial u}{\partial x} \frac{\partial v}{\partial z} - F_2(y, u) \frac{\partial^2 u}{\partial x \partial z} \frac{\partial u}{\partial z} \frac{\partial v}{\partial y}}{\frac{\partial u}{\partial x} \frac{\partial v}{\partial y}} \quad . \end{aligned} \quad (18)$$

By inserting equation (18) into equation (14), we get

$$\begin{aligned} F_2(x, v) \frac{\partial^2 v}{\partial z^2} - F_2(y, u) \frac{\partial^2 u}{\partial z^2} &= \frac{F_2(x, v) \frac{\partial^2 v}{\partial z \partial y} \frac{\partial u}{\partial x} \frac{\partial v}{\partial z} - F_2(y, u) \frac{\partial^2 u}{\partial x \partial z} \frac{\partial u}{\partial z} \frac{\partial v}{\partial y}}{\frac{\partial u}{\partial x} \frac{\partial v}{\partial y}} \\ \Leftrightarrow \frac{F_2(x, v)}{F_2(y, u)} \left(\frac{\partial^2 v}{\partial z^2} \frac{\partial u}{\partial x} \frac{\partial v}{\partial y} - \frac{\partial^2 v}{\partial z \partial y} \frac{\partial u}{\partial x} \frac{\partial v}{\partial z} \right) &= \frac{\partial^2 u}{\partial z^2} \frac{\partial u}{\partial x} \frac{\partial v}{\partial y} - \frac{\partial^2 u}{\partial z \partial x} \frac{\partial u}{\partial z} \frac{\partial v}{\partial y} . \end{aligned} \quad (19)$$

Notice that by equation (11) we can write $\frac{F_2(x, v)}{F_2(y, u)} = \frac{\partial u}{\partial z} / \frac{\partial v}{\partial z}$. Inserting this into equation (19), and dividing by $\frac{\partial v}{\partial y} \frac{\partial u}{\partial x} \frac{\partial u}{\partial z}$, the equation simplifies to

$$\frac{\frac{\partial^2 v}{\partial z^2}}{\frac{\partial v}{\partial z}} - \frac{\frac{\partial^2 v}{\partial z \partial y}}{\frac{\partial v}{\partial y}} = \frac{\frac{\partial^2 u}{\partial z^2}}{\frac{\partial u}{\partial z}} - \frac{\frac{\partial^2 u}{\partial z \partial x}}{\frac{\partial u}{\partial x}} .$$

The equation can be written also as

$$\begin{aligned} \frac{\partial}{\partial z} \ln \frac{\partial v}{\partial z} - \frac{\partial}{\partial z} \ln \frac{\partial v}{\partial y} &= \frac{\partial}{\partial z} \ln \frac{\partial u}{\partial z} - \frac{\partial}{\partial z} \ln \frac{\partial u}{\partial x} \\ \Leftrightarrow \frac{\partial}{\partial z} \ln \left(\frac{\frac{\partial v}{\partial z}}{\frac{\partial v}{\partial y}} \right) &= \frac{\partial}{\partial z} \ln \left(\frac{\frac{\partial u}{\partial z}}{\frac{\partial u}{\partial x}} \right) . \end{aligned} \quad (20)$$

Since the left hand side depends on y, z and right hand side depends on x, z , it follows that both must be functions of only z . Furthermore, since the derivative of the logarithm of a function is a function of z , the function itself must be of the form

$$\frac{\frac{\partial v}{\partial z}}{\frac{\partial v}{\partial y}} = \frac{\Phi_1(y)}{\Phi_2(z)} . \quad (21)$$

On the other hand, dividing equation (12) by equation (13), we get

$$\frac{F_2(x, v)}{F_1(x, v)} \frac{\partial v}{\partial y} = \left(\frac{F_2(y, u)}{F_1(y, u)} \frac{\partial u}{\partial x} \right)^{-1} . \quad (22)$$

By equation (21), $\frac{F_2(x, v)}{F_1(x, v)} = \frac{\Phi_1(x)}{\Phi_2(v)}$. Inserting this, we get

$$\begin{aligned} \frac{\Phi_1(x)}{\Phi_2(v)} \frac{\partial v}{\partial y} &= \left(\frac{\Phi_1(y)}{\Phi_2(u)} \frac{\partial u}{\partial x} \right)^{-1} \\ \Leftrightarrow \frac{\Phi_1(y)}{\Phi_2(v)} \frac{\partial v}{\partial y} &= \left(\frac{\Phi_1(x)}{\Phi_2(u)} \frac{\partial u}{\partial x} \right)^{-1} . \end{aligned} \quad (23)$$

Now, since left hand side depends only on (y, z) and right hand side on (x, z) , each must be a function of z only, that is $g(z)$. Furthermore, we note that we must have

$$g(z) = \frac{1}{g(z)} \Leftrightarrow g(z) = \pm 1 . \quad (24)$$

Since the condition must be fulfilled for all x, y , we must have

$$\begin{aligned}\frac{\Phi_1(y)}{\Phi_2(v)} \frac{\partial v}{\partial y} &= 1 \\ \left(\frac{\Phi_1(y)}{\Phi_2(v)} \frac{\partial v}{\partial y} \right)^{-1} &= 1\end{aligned}$$

for each y as well. Summing these, we get

$$\begin{aligned}\frac{\Phi_1(y)}{\Phi_2(v)} \frac{\partial v}{\partial y} + \frac{\Phi_2(v)}{\Phi_1(y)} \frac{1}{\frac{\partial v}{\partial y}} &= 2 \\ \Leftrightarrow \left(\frac{\partial v}{\partial y} - \frac{\Phi_2(v)}{\Phi_1(y)} \right)^2 &= 0 \quad .\end{aligned}$$

We can then write

$$\begin{aligned}\frac{\partial v}{\partial y} &= \frac{\Phi_2(v)}{\Phi_1(y)} \\ \frac{\partial v}{\partial z} &= \frac{\Phi_2(v)}{\Phi_2(z)} \quad .\end{aligned}$$

Combining these to obtain the differential dv we get

$$\frac{dv}{\Phi_2(v)} = \frac{dy}{\Phi_1(y)} + \frac{dz}{\Phi_2(z)} \quad . \quad (25)$$

By denoting $\int \frac{dp}{\Phi_k(p)} = \ln f_k(p)$, we obtain

$$C f_2(v) = f_1(y) f_2(z) \quad . \quad (26)$$

The function f_1 can be incorporated into our model, that is $f_1 \circ h \mapsto h$. By inserting $v = F(y, z)$ we get the final form

$$C f_2(F(y, z)) = h(y) f_2(z) \quad . \quad (27)$$

□

A.2 Mapping from $p(c \mid \mathbf{x}, \theta)$ to $h(\mathbf{r}, \theta)$ is Monotonically Increasing

Proposition A.1 *From axiom 5 it follows that*

$$\log h(\mathbf{r}, \theta) = f_C(\log p(c \mid \mathbf{x}, \theta)) \quad (28)$$

where f_C is a monotonically increasing function.

Proof Denoting $\tilde{\theta}_r = p(\mathbf{r} \mid \tilde{\theta})$ in inequalities (5) and (6) in the paper we can write them in the following form

$$\sum_{\mathbf{r} \in R} \tilde{\theta}_r \log h(\mathbf{r}, \theta_1) \leq \sum_{\mathbf{r} \in R} \tilde{\theta}_r \log h(\mathbf{r}, \theta_2) \quad (29)$$

$$\begin{aligned} & \Updownarrow \\ & \sum_{\mathbf{r} \in R} \tilde{\theta}_r \log p(c|\mathbf{x}, \theta_1) \leq \sum_{\mathbf{r} \in R} \tilde{\theta}_r \log p(c|\mathbf{x}, \theta_2) \quad . \end{aligned} \quad (30)$$

Consider the points in the parameter space Θ , where $\tilde{\theta}_k = 1$ and $\tilde{\theta}_i = 0$ for $i \neq k$ (“corner points”). In these points the linear combinations vanish and the equivalent inequalities (29) and (30) become

$$\left\{ \begin{array}{l} \log h(\mathbf{r}_k, \theta_1) \leq \log h(\mathbf{r}_k, \theta_2) \\ \Updownarrow \\ \log p(c_k|\mathbf{x}_k, \theta_1) \leq \log p(c_k|\mathbf{x}_k, \theta_2) \end{array} \right. . \quad (31)$$

Since the functional form of f_C must be the same regardless of the choice of $\tilde{\theta}$, equivalence (31) holds everywhere in the parameter space, not just in the corners.

From the equivalence (31) (and the symmetry of the models with respect to re-labeling the data items) it follows that $h(\mathbf{r}, \theta)$ must be of the form

$$\log h(\mathbf{r}, \theta) = f_C(\log p(c|\mathbf{x}, \theta))$$

where f_C is a monotonically increasing function. \square

A.3 Mapping is of Form $h(\mathbf{r}, \theta) = \exp(\beta) p(c | \mathbf{x}, \theta)^A$

Proposition A.2 *For a continuous increasing function $f_C(t)$ for which*

$$\log h(\mathbf{r}, \theta) = f_C(\log p(c|\mathbf{x}, \theta))$$

it follows from axiom 5 that

$$f_C(t) = At + \beta \quad ,$$

or, equivalently,

$$h(\mathbf{r}, \theta) = \exp(\beta) p(c|\mathbf{x}, \theta)^A \quad \text{with} \quad A > 0. \quad (32)$$

Proof Note, that we can decompose K_{COND} as

$$K_{COND}(\tilde{\theta}, \theta) = S(\tilde{\theta}) - R(\tilde{\theta}, \theta) \quad , \quad (33)$$

where $S(\tilde{\theta}) = \sum_{p(|\tilde{\theta})} \log p(c|\mathbf{x}\tilde{\theta})$ and $R(\tilde{\theta}, \theta) = \sum_{p(|\tilde{\theta})} \log p(c|\mathbf{x}\theta)$. Consider any $\tilde{\theta}$ and the set of points θ that satisfy $R(\tilde{\theta}, \theta) = t$, where t is some constant. From the fifth axiom (the equality part) it follows that there must exist a constant $f_{\tilde{\theta}}(t)$ that defines the same set of points θ , defined by $\sum_{\mathbf{r} \in R} p(\mathbf{r}|\tilde{\theta}) \log h(\mathbf{r}, \theta) = f_{\tilde{\theta}}(t)$. From

the inequality part of the same axiom it follows that $f_{\tilde{\theta}}$ is a monotonically increasing function. Hence,

$$\sum_{\mathbf{r} \in R} p(\mathbf{r}|\tilde{\theta}) \log h(\mathbf{r}, \theta) = f_{\tilde{\theta}} \left(\sum_{\mathbf{r} \in R} p(\mathbf{r}|\tilde{\theta}) \log p(c|\mathbf{x}, \theta) \right) . \quad (34)$$

On the other hand, from equation (28) we know that we can write

$$\log h(\mathbf{r}, \theta) = f_C (\log p(c|\mathbf{x}, \theta)) . \quad (35)$$

So, equations (34) and (35) lead to

$$f_{\tilde{\theta}} \left(\sum_{\mathbf{r} \in R} p(\mathbf{r}|\tilde{\theta}) \log p(c|\mathbf{x}, \theta) \right) = \sum_{\mathbf{r} \in R} p(\mathbf{r}|\tilde{\theta}) f_C (\log p(c|\mathbf{x}, \theta)) . \quad (36)$$

If we make a variable change $u_i = \log p(c_i | \mathbf{x}_i, \theta)$ and denote $p(\mathbf{r}_i|\tilde{\theta}) = \tilde{\theta}_i$ for brevity, equation (36) becomes

$$f_{\tilde{\theta}} \left(\sum_i \tilde{\theta}_i u_i \right) = \sum_i \tilde{\theta}_i f_C (u_i) . \quad (37)$$

Not all u_i are independent, however: for each fixed \mathbf{x} , one of the variables u_l is determined by the other u_i 's

$$\exp(u_i) = p(c_i | \mathbf{x}_i, \theta) ,$$

and

$$\sum_{\text{fixed } \mathbf{x}} p(c_i | \mathbf{x}, \theta) = 1 \Leftrightarrow \sum_{\text{fixed } \mathbf{x}} \exp(u_i) = 1 .$$

So the last u_l for each \mathbf{x} is

$$u_l = \log \left(1 - \sum_{\substack{\text{fixed } \mathbf{x} \\ \text{indep. } u_m}} \exp(u_m) \right) , \quad (38)$$

where the sum only includes the independent variables u_m for the fixed \mathbf{x} . This way we can make the dependency on each u_i explicit in equation (37):

$$\begin{aligned} & f_{\tilde{\theta}} \left[\sum_{\text{indep. } u_j} \tilde{\theta}_j u_j + \sum_{\text{dependent } u_l} \tilde{\theta}_l \log \left(1 - \sum_{\substack{\text{fixed } \mathbf{x} \\ \text{indep. } u_m}} \exp(u_m) \right) \right] \\ &= \sum_{\text{indep. } u_j} \tilde{\theta}_j f_C(u_j) + \sum_{\text{dependent } u_l} \tilde{\theta}_l f_C \left(\log \left(1 - \sum_{\substack{\text{fixed } \mathbf{x} \\ \text{indep. } u_m}} \exp(u_m) \right) \right) . \end{aligned} \quad (39)$$

Let us differentiate both sides with respect to a u_k :

$$\begin{aligned}
& \underbrace{f'_{\tilde{\theta}} \left(\sum_i \tilde{\theta}_i u_i \right)}_{\alpha} \left[\tilde{\theta}_k - \underbrace{\frac{\tilde{\theta}_l}{u_l}}_{c_x} \exp(u_k) \right] \\
&= \tilde{\theta}_k f'_C(u_k) - \underbrace{\frac{\tilde{\theta}_l}{u_l}}_{c_x} \underbrace{f'_C(u_l)}_{d_x} \exp(u_k) \quad . \tag{40}
\end{aligned}$$

For all such variables u_k that share the same \mathbf{x} , we get

$$\begin{aligned}
\alpha \left[\tilde{\theta}_k - c_x \exp(u_k) \right] &= \tilde{\theta}_k f'_C(u_k) - c_x d_x \exp(u_k) \\
&\Updownarrow \\
\tilde{\theta}_k \exp(-u_k) (f'_C(u_k) - \alpha) &= c_x (d_x - \alpha) \quad .
\end{aligned}$$

Since the right-hand side only depends on \mathbf{x} , not on individual u_k , the left-hand side must also only depend on \mathbf{x} and the factors depending on u_k must cancel out.

$$\begin{aligned}
f'_C(u_k) - \alpha &= B_x \frac{\exp(u_k)}{\tilde{\theta}_k} \\
&\Updownarrow \\
\underbrace{f'_{\tilde{\theta}} \left(\sum_i \tilde{\theta}_i u_i \right)}_{\text{does not depend on } u_k} &= \underbrace{f'_C(u_k) - B_x \frac{\exp(u_k)}{\tilde{\theta}_k}}_{\text{depends on } u_k} \quad .
\end{aligned}$$

Since the left-hand side depends neither on u_k nor \mathbf{x} , both sides must be constant

$$\begin{aligned}
&\implies f'_{\tilde{\theta}}(t) = A \\
&\implies f_{\tilde{\theta}}(t) = A t + \beta \quad . \tag{41}
\end{aligned}$$

Substituting (41) into equation (37) we get

$$\begin{aligned}
A \left(\sum_i \tilde{\theta}_i u_i \right) + \beta &= \sum_i \tilde{\theta}_i f_C(u_i) \\
&\Updownarrow \\
\sum_i \tilde{\theta}_i (A u_i - f_C(u_i)) &= -\beta \quad ,
\end{aligned}$$

and since this must hold for any parameters $\tilde{\theta}$, it must also hold for the corner points:

$$\begin{aligned}
A u_i - f_C(u_i) &= -\beta \\
\implies f_C(t) &= A t + \beta \quad . \tag{42}
\end{aligned}$$

□

A.4 Axiom 6 Implies Exponent $A = 1$

Proposition A.3 *From axiom 6 it follows that $A = 1$.*

Proof Without axiom 6, the discriminative posterior would be unique up to a positive constant A :

$$p_d(\theta | D) \propto p(\theta) \prod_{(c, \mathbf{x}) \in D} p(c | \mathbf{x}, \theta)^A. \quad (43)$$

Axiom 6 is used to fix this constant to unity by requiring the discriminative posterior to obey the Bayesian convention for a fixed \mathbf{x} , that is, the discriminative modeling should reduce to Bayesian joint modeling when there is only one covariate. We require that for a fixed \mathbf{x} and a data set $D_x = \{(c, \mathbf{x}') \in D \mid \mathbf{x}' = \mathbf{x}\}$ the discriminative posterior matches the joint posterior p_j^x of a model p^x , where $p^x(c | \theta) \equiv p(c | \mathbf{x}, \theta)$,

$$p_j^x(\theta | D_x) \propto p(\theta) \prod_{c \in D_x} p^x(c | \theta) \propto p(\theta) \prod_{(c, \mathbf{x}) \in D_x} p(c | \mathbf{x}, \theta), \quad (44)$$

Clearly, $A = 1$ satisfies the axiom 6, i.e., $p_j^x(\theta | D_x)$ equals $p_d(\theta | D_x)$ for all $\theta \in \bar{\Theta}$. If the proposal would be false, the axiom 6 should be satisfied for some $A \neq 1$ and for all θ and data sets. In particular, the result should hold for a data set having a single element, (c, \mathbf{x}) . The discriminative posterior would in this case read

$$p_d(\theta | D_x) = \frac{1}{Z_A} p(\theta) p(c | \mathbf{x}, \theta)^A, \quad (45)$$

where $D_x = \{(c, \mathbf{x})\}$ and Z_A is a normalization factor, chosen so that the posterior satisfies $\int p_d(\theta | D) d\theta = 1$. The joint posterior would, on the other hand, read

$$p_j^x(\theta | D_x) = \frac{1}{Z_1} p(\theta) p(c | \mathbf{x}, \theta). \quad (46)$$

These two posteriors should be equal for all θ :

$$1 = \frac{p_d(\theta | D_x)}{p_j^x(\theta | D_x)} = \frac{Z_1}{Z_A} p(c | \mathbf{x}, \theta)^{A-1}. \quad (47)$$

Because the normalization factors Z_1 and Z_A in equation (47) are constant in θ , also $p(c | \mathbf{x}, \theta)^{A-1}$ must be constant in θ . This is possible only if $A = 1$ or $p(c | \mathbf{x}, \theta)$ is a trivial function (constant in θ) for all c .

□